AD-A273 638

High-Level Connectionist Models

Jordan B. Pollack
Laboratory for Artificial Intelligence Research
Department of Computer and Information Science

S DTIC
ELECTE
DEC 13 1993
A

93-29857

October 1993

93 12 7 053

REPORT DOCUMENTATION PAGE

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	October 1993	Final

4. TITLE AND SUBTITLE

High-Level Connectionist Models

5. FUNDING NUMBERS

N00014-93-1-0059

6. AUTHOR(S)

Jordan B. Pollack

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

The Ohio State University Research Foundation
1960 Kenny Rd.
Columbus, Ohio 43210

8. PERFORMING ORGANIZATION REPORT NUMBER

RF No. 760388/726944

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Department of the Navy
Office of Naval Research
800 N. Quincy St.
Arlington, Virginia 22217-5660

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

See report

14. SUBJECT TERMS

15. NUMBER OF PAGES

44

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT

High-Level Connectionist Models

Jordan B. Pollack
Laboratory for Artificial Intelligence Research
Department of Computer and Information Science

Accesion For		
NTIS CRA&I		N
DTIC TAB		[]
Unannounced		[]
Justification		
By		
Distribution /		
Availability Codes		
Dist	Avail and / or Special	
A-1		

October 1993

Jordan B. Pollack
Ohio State University
(614)-292-4890
pollack@cis.ohio-state.edu
High-Level Connectionist Models
N00014-93-1-0059
1 Oct 92- 30 Sep 93

1.0 Productivity Measures

2	Refereed papers submitted
2	Refereed papers published
0	Unrefereed Reports
0	Books of parts thereof submitted
0	Books or parts thereof published
0	Patents filed but not granted
0	Patents granted
2	Invited Presentations
2	Contributed Presentations
1	Honors
0	Prizes
0	Promotions
1	Graduate Students Supported
0	Minorities Supported

Jordan B. Pollack
Ohio State University
(614)-292-4890
pollack@cis.ohio-state.edu
High-Level Connectionist Models
N00014-93-1-0059
1 Oct 92- 30 Sep 93

2.0 Summary of Technical Results

Our goals this year involved learning in connectionist networks while automatically decomposing behaviors in order to support those behaviors with modular architectures. While there has been some work in this area, (e.g. Jacobs, Jordan & Barto 1991) we desired to have the modules fully evolve in response to the demands of the task. To accomplish this, we needed a training mechanism more robust than backpropagation, so we turned towards *genetic algorithms* (GAs). These algorithms, based on principles adopted from natural selection, allow solutions to be evolved which fit the requirements of an environment.

There is an extensive body of work applying GAs to evolving neural networks, but most simply use GAs to set the weights for a fixed-structure network. Those that attempt to evolve network structure do so in a very limited way. (See Schaffer, et al., 1992 for a good overview.) Thus before applying GAs to network modularization, we first had to solve the "generalized network acquisition" problem, i.e., the problem of acquiring both network structure and weight values simultaneously. The result was GNARL, an algorithm for *GeNeralized Acquisition of Recurrent Links* (Angeline, Saunders, and Pollack, 1993). In contrast to constructive and destructive algorithms for network induction (e.g., Fahlman and Lebiere, 1990; Fahlman 1991; Chen, et al., 1993; Mozer & Smolensky, 1989; Omlin & Giles, 1993), GNARL is an "unsupervised" form of network learning using a population of mutating networks where the severity of mutation is modulated by feedback from an fitness function - as a network approaches the goal, its chances of reproducing without modification decrease towards 0.

The power of GNARL has been demonstrated on several tasks, including the Tomita language acquisition task (Pollack, 1991). Here we will highlight the *Tracker* task, described by Jefferson, et al. (1991). In this problem, a simple agent (a simulated ant) is placed on a two-dimensional toroidal grid that contains a trail of food. The ant traverses the grid, eating in one time step any food it contacts. The goal of the task is to maximize the number of pieces of food the ant eats within a predefined allotted time. The trail of food used in the experiment (shown in figure

1a) was hand-crafted in the original study to be increasingly difficult for the evolving ants.


(a)

(b)

(c)

Figure 1: The ant problem. (a) The trail is connected initially, but becomes progressively more difficult to follow. The underlying 2-d grid is toroidal, so that position "P" is the first break in the trail. (b) The semantics of the I/O units for the ant network. The first input node denotes the presence of food in the square directly in front of the ant; the second denotes the absence of food in this same square. No-op, from Jefferson, allows the network to stay in one position while activation flows through recurrent links. The network in (b)"eats" 42 pieces of food before spinning endlessly in place at position P. The simple FSA in (c) is a handcrafted solution which moves if there is food, otherwise turns and looks in 4 directions before continuing to move.

Following Jefferson, *et al* (1991), the ant is controlled by a network with two input nodes and four output nodes, as shown in figure 1b. The first input node denotes the presence of food in the square directly in front of the ant; the second denotes the absence of food in this same square, restricting the possible legal inputs to the network to (1, 0) or (0, 1). Each of the four output units corresponds to a unique executable action – move forward, turn left, turn right, or no-op. Each ant is in an implicit sense/act loop that repeatedly sets the input activations of the network, computes the activations of the output nodes, determines the output node with

maximum activation and executes its associated action. Every application of the sense/act loop is assumed to happen in a single time step. Once a position with food is visited, the food is removed. The fitness function used in this task is simply the number of food units consumed in 200 time steps.

In these experiments, we used a population of 100 networks. In the first run (2090 generations using 104,600 network evaluations), GNARL found a network that cleared 81 grid positions within the 200 time steps. When we allowed this ant to run for an additional 119 time steps, it successfully cleared the entire trail.

Analysis of the behavior of the network under conditions of constant food inputs and constant no-food inputs revealed that it had the same behavior as a simple 5-state automata designed by Jefferson(figure 1c). In an empty torus, it would simply move across a line forever. GNARL was not constrained to find FSA's however. On another run it evolved a network that cleared 82 grid points within the 200 time steps. Figure 3 shows the limit behavior of this network under different input conditions. Each graph shows a collection of the output units as a point in the 3D space (move,right,left), since no-op is never used. Under the condition of constant food (a), it moves. In the case of a no-food condition, instead of spinning in one place, or rotating 4 times and moving, this network makes a complicated pseudo-random series of turns and moves, which shows up as the "D" shape in the graph, and enables it to explore more of the empty torus. Figure 3(c) shows the output of the network on the given trail, and figure 3(d) plots the movement of the agent over time on an empty grid.

2.1 Cooperating Agents

A natural next question to ask is whether such simple evolved agents can behave in a cooperative manner. The issue of evolved cooperation has been the subject of intense research since Axelrod (1984) showed that two agents engaged in a prisoner's dilemma (Poundstone, 1992) can evolve into mutually cooperative states. Recent studies explore alternatives to Axelrod's Tit-for-Tat strategy (Nowak and Sigmund, 1993); other research extends the two-way interactions by introducing multiple agents. Such work is still far from complete – some agents that are completely designed (e.g., Drogoul and Ferber, 1993); some have homogenous structure and perform only parametric learning (e.g., Assad and Packard, 1992); and finally some focus on simple tasks (e.g., Deneubourg, et al., 1991).

Our most recent work is the study of how to evolve cooperation among heterogenous agents. *Our thesis is that a set of such individuals – associated only loosely by being part of the same group, and communicating only indirectly via the environment – should be able to learn to cooperate to solve a given task.*

Figure 2: Limit behavior of a network for the ant problem. Graphs show the state of the output units Move, Right, Left. (a) Fixed point attractor that results for sequence of 500 "food" signals; (b) Limit cycle attractor that results when a sequence of 500 "no food" signals is given to network; (c) All states visited while traversing the trail; (d) The x position of the ant over time when run on an empty grid.

Our agents are controlled by recurrent neural networks, and evolved using GNARL. Because GNARL places no restrictions on the number of hidden nodes of a network or on the network's connectivity, it is an ideal system in which to build cooperative agents with heterogenous structure.

Motivated by biology, many researchers have expanded upon the ant-domain to explore cooperative problem-solving in insect-like agents (e.g., Colorni, et al., 1992; Drogoul, et al., 1992; Deneubourg et al., 1991; Theraulaz, et al., 1991). We

have chosen an alternate domain, one which we feel should serve as a rich substrate for our theories of cooperative evolution: soccer.

A competitive team sport like soccer is a complex task. Cooperation is much more direct than in insect "food-collection" tasks which can often be seen as simple questions of parallelism: players must coordinate defensively to protect the goal by covering different opposition team members, and must coordinate offensively to score. Passing – the most critical aspect of the game – cannot be accomplished unless both the kicker and the recipient work closely together. Additionally, there is the added complication of competition between opposing teams.

The dual elements of cooperation and competition make credit assignment in soccer particularly difficult. Consider evaluating a given agent. Even if it is the best player on the field, its team may lose. Likewise, if it is the worst player on the field, its team may win. Moreover, the game may result in a tie. How is the fitness of each team (player) evaluated in this case? Taking motivation from the principles of soccer, we are currently exploring two key factors: ball possession (the amount of time a team controls the ball) and support (the amount of aid a team gives an agent who controls the ball).

At this stage in our work, we have developed a set of agent sensors and actuators that is sufficient for cooperation in soccer, shown in Tables 1 and 2.

As a proof of concept, we have both designed by hand, and evolved by network, several soccer players using these primitives. To demonstrate, a set of snapshots of a game is shown in Figure 3. Play begins with the *Circle* team, defending the lower goal, and the *Square* team, defending the upper (Figure 3, top-left). Members of both teams execute move-to-ball (top-right), but because players can are heterogenous, player 4 of the *Square* team is fastest and reaches the ball first (bottom-left). Sensors left-open, forward-open, and right-open return negative results, so player 4 passes back to its teammate, player 7 (bottom-right). Beyond the scenes depicted in the figure, player 7 passes back to player 6, who then passes up to player 4, who dribbles forward but eventually loses control of the ball, demonstrating that our chosen primitives support both cooperation and competition

There are several reasons for the relatively high level of modeling we have chosen. First, to explore cooperation in soccer, something like our sensors/ behaviors must exist. Had we chosen a lower-level set of primitives, concepts such as "goal-open" would still have had to emerge. The situation is analogous to the instructions for coaches in introductory soccer books: they begin with "how to kick the ball," then progress to "how to pass", and maybe in the last chapter

Table 1: Soccer Sensors

Sensor	Meaning
touching-ball	player is touching ball
goal-open	player has an open shot at goal
left-open	area to left of player is clear
forward-open	area in front of player is clear
right-open	area to right of player is clear
open-teammate-p	player has a teammate open
better-teammate-p	player has a teammate open who is "better" (closer to goal or more open)
we-possess	our team has the ball

Table 2: Soccer Actuators

Actuators	Meaning
move-to-ball	move one step in direction of ball
shoot-to-goal	kick the ball hard to the goal
dribble-forward	kick the ball gently forward
dribble-left	kick the ball gently to the left
dribble-right	kick the ball gently to the right
pass	pass the ball to the best open teammate
get-open	move to an open position
try-to-trap	stop the ball from moving

quickly discuss team-level strategies (e.g., Hargreaves, 1990; McGettian, 1987). Since we are mostly interested in team-level issues, we chose a high-set of primitives for our starting point. Secondly, the starting level of abstraction is largely irrelevant – what is important is the distance between the starting primitives and the ultimate behavior. Finally, the level we have chosen is still too complex to completely solve the task by hand. We have already modeled soccer players with these primitives, and even designed expert players, but in preliminary studies, the evolved networks often outperformed the hand built experts. It is difficult to predict how heterogenous players will work together.

Distancing ourselves from this particular game domain, future work should provide deeper insights into the artificial evolution of cooperation and complexity. How should the fitness of a group of agents acting cooperatively be evaluated? Is it necessary to assign fitness to individual players, or will group scoring be enough. Either less able players will luckily be weeded out by the competitive process of teams, or they will hitchhike along. To evolve cooperating

Figure 3: Snapshots of passing. Top-left: start of game. Top-right: movement to ball. Bottom-left: Player 4 poised to pass. Bottom-right: player 4 has passed back to player 7.

agents, must each player always cooperate, or is there enough flexibility to allow for the selfish behavior found in groups (e.g. show-off players)? What are the dynamics in such open games which create and "lock in" evolutionarily stable but mediocre strategies, and how can that lock be broken? We are now poised to ask these and other questions; a team competition as modelled by our soccer simulation is a fertile ground for research into cooperative and competitive teams of heterogenous agents.

2.2 References

Angeline, P. J. and Pollack, J. B. (1993). Competitive environments evolve better solutions for complex tasks. In Forrest, S., editor, *Genetic Algorithms: Proceedings of the Fifth International Conference*.

Angeline, P. J., Saunders, G. M., and Pollack, J. B. (1994). An evolutionary algorithm that constructs recurrent neural knotweeds. *IEEE Transactions on Neural Networks*, to appear.

Assad, A. M. and Packard, N. H. (1992). Emergent colonization in an artificial colony. In Varela, F. J. and Bourgine, P., editors, *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 143–152. MIT Press.

Axelrod, R. (1984). *The Evolution of Cooperation*. Basic Books.

Chen, D., Giles, C., Sun, G., Chen, H., Less, Y., and Goudreau, M. (1993). Constructive learning of recurrent neural networks. *IEEE International Conference on Neural Networks*, 3:1196–1201.

Colorni, A., Dorigo, M., and Maniezzo, V. (1992). Distributed optimization by ant colonies. In Varela, F. J. and Bourgine, P., editors, *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 143–152. Francisco J. Varela and Paul Bourgine.

Deneubourg, J. L., Goss, S., Franks, N., Sendova-Franks, A., Detrain, C., and Chrétien, L. (1991). The dynamics of collective sorting: Robot-like ants and ant-like robots. In Meyer, J. A. and Wilson, S. W., editors, *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, pages 356–363. MIT Press.

Drogoul, A. and Ferber, J. (1993). From tom thumb to the dockers: Some experiments with foraging robots. In Meyer, J. A., Roitblat, H. L., and Wilson, S. W., editors, *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, pages 451–459. MIT Press.

Drogoul, A., Ferber, J., Corbara, B., and Fresneau, D. (1992). A behavioral simulation model for the study of emergent social structures. In Varela, F. J. and Bourgine, P., editors, *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 161–170. MIT Press.

Fahlman, S. (1991). The recurrent cascade-correlation architecture. In Lippmann, R., Moody, J., and Touretzky, D., editors, *Advances in Neural Information Processing Systems 3*, pages 190–196. Morgan Kaufmann, San Mateo.

Fahlman, S. E. and Lebiere, C. (1990). The cascade-correlation architecture. In Touretzky, D. S., editor, *Advances in Neural Information Processing Structures 2*, pages 524–532. Morgan Kaufmann.

Greenfield, M. D. and Rolzen, I. (1993). Katydid synchronous chorusing is an evolutionarily stable outcome of female choice. *Nature*, 364:618–620.

Hargreaves, A. (1990). *Skills and Strategies for Coaching Soccer*. Leisure Press, Campaign, Illinois.

Jefferson, D., Collins, R., Cooper, C., Dyer, M., Flowers, M., Korf, R., Taylor, C., and Wang, A. (1991). Evolution as a theme in artificial life: The Genesys/Tracker system. In Langton, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II: Proceedings of the Workshop on Artificial Life*, pages 549–577. Addison-Wesley.

Koza, J. (1992). Genetic evolution and co-evolution of computer programs. In Christopher G. Langton, Charles Taylor, J. D. F. and Rasmussen, S., editors, *Artificial Life II*. Addison Wesley Publishing Company, Reading Mass.

McGettigan, J. P. (1987). *Soccer drills for team and individual play*. Prentice Hall.

Mozer, M. and Smolensky, P. (1989). Skeletonization: A technique for trimming the fat from a network via relevance assessment. In Touretzky, D., editor, *Advances in Neural Information Processing Systems 1*, pages 107–115. Morgan Kaufmann, San Mateo.

Nowak, M. and Sigmund, K. (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature*, 364:56–58.

Omlin, C. W. and Giles, C. L. (1993). Pruning recurrent neural networks for improved generalization performance. Technical Report Tech Raport No 93-6, Computer Science Department, Rensselaer Polytechnic Institute.

Poundstone, W. (1992). *Prisoner's Dilemma*. Doubleday.

Saunders, G. M., Angeline, P. J., and Pollack, J. B. (1994). Structural and behavioral evolution of recurrent networks. In *Advances in Neural Information Processing 7*, to appear. Morgan Kaufmann.

Saunders, G. M., Kolen, J. F., Angeline, P. J., and Pollack, J. B. (1992). Additive modular learning in preemptrons. In *Proceedings of the 14th Annual Conference of the Cognitive Science Society*, pages 1098–1103, Bloomington, Indiana.

Theraulaz, G., Goss, S., Gervet, J., and Deneubourg, J. L. (1991). Task differentiation in Polistes wasp colonies: A model for self-organizing groups of robots. In Meyer, J. A. and Wilson, S. W., editors, *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, pages 346–355. MIT Press.

Jordan B. Pollack
Ohio State University
(614)-292-4890
pollack@cis.ohio-state.edu
High-Level Connectionist Models
N00014-93-1-0059
1 Oct 92- 30 Sep 93

3.0 Publication Activity

3.1 Refereed Papers

1. Angeline, P., Saunders, G., & Pollack, J. (in Press) An evolutionary Algorithm that constructs recurrent neural networks. IEEE Trans. on Neural Nets

2. Saunders, G. Angeline, P. & Pollack, J. (Accepted) Structural and Behavioral Evolution of recurrent networks.Neural Information Processing Systems Conference,

3.2 Presentations on this work

1. Saunders presented work at the Midwest Connectfest III, held November 1992 at Carnegie-Mellon.

2. Angeline & Saunders presented this work in workshops at the International Conference on Genetic Algorithms, July 1993, at University of Illinois.

3. Pollack was external faculty at Santa Fe Institute in Nov. 1992, and presented some of this work.

4. Pollack was keynote speaker "Only 99 (hundred million) lines of code to go" at the Midwest Artificial Intelligence and Cognitive Science meeting, April 1993 in Indiana Dunes.

Jordan B. Pollack
Ohio State University
(614)-292-4890
pollack@cis.ohio-state.edu
High-Level Connectionist Models
N00014-93-1-0059
1 Oct 92- 30 Sep 93

4.0 Transitions & Dept. of Defense Interaction

None

Jordan B. Pollack
Ohio State University
(614)-292-4890
pollack@cis.ohio-state.edu
High-Level Connectionist Models
N00014-93-1-0059
1 Oct 92- 30 Sep 93

5.0 Software Prototypes

We have built and released an efficient set of software tools for research in recurrent neural nets. For ease of construction, we implemented the system in Common Lisp for its flexibility. But in order to make the resultant system more efficient, we circumscribed all numerically intensive subroutines in a small well-defined library of linear algebra routines. This has been implemented both in LISP and in C, but the foreign function interface is specific to the foreign function capacity of the "Allegro" brand of Common Lisp, and is not completely portable. We have supplemented these libraries with new modes of learning based on genetic algorithms and genetic programming.

The lisp software is available to interested researchers by FTP free of charge. Please contact ansil@cis.ohio-state.edu for further information.

Appendix

An Evolutionary Algorithm that Constructs Recurrent Neural Networks

Peter J. Angeline, Gregory M. Saunders and Jordan B. Pollack

Laboratory for Artificial Intelligence Research
Computer and Information Science Department
The Ohio State University
Columbus, Ohio 43210
pja@cis.ohio-state.edu
saunders@cis.ohio-state.edu
pollack@cis.ohio-state.edu

Abstract

Standard methods for inducing both the structure and weight values of recurrent neural networks fit an assumed class of architectures to every task. This simplification is necessary because the interactions between network structure and function are not well understood. Evolutionary computation, which includes genetic algorithms and evolutionary programming, is a population-based search method that has shown promise in such complex tasks. This paper argues that genetic algorithms are inappropriate for network acquisition and describes an evolutionary program, called GNARL, that simultaneously acquires both the structure and weights for recurrent networks. This algorithm's empirical acquisition method allows for the emergence of complex behaviors and topologies that are potentially excluded by the artificial architectural constraints imposed in standard network induction methods.

An Evolutionary Algorithm that Constructs Recurrent Neural Networks

Peter J. Angeline, Gregory M. Saunders and Jordan B. Pollack

Laboratory for Artificial Intelligence Research
Computer and Information Science Department
The Ohio State University
Columbus, Ohio 43210
pja@cis.ohio-state.edu
saunders@cis.ohio-state.edu
pollack@cis.ohio-state.edu

Abstract

Standard methods for inducing both the structure and weight values of recurrent neural networks fit an assumed class of architectures to every task. This simplification is necessary because the interactions between network structure and function are not well understood. Evolutionary computation, which includes genetic algorithms and evolutionary programming, is a population-based search method that has shown promise in such complex tasks. This paper argues that genetic algorithms are inappropriate for network acquisition and describes an evolutionary program, called GNARL, that simultaneously acquires both the structure and weights for recurrent networks. This algorithm's empirical acquisition method allows for the emergence of complex behaviors and topologies that are potentially excluded by the artificial architectural constraints imposed in standard network induction methods.

1.0 Introduction

In its complete form, network induction entails both *parametric* and *structural* learning [1], i.e., learning both weight values and an appropriate topology of nodes and links. Current methods to solve this task fall into two broad categories. *Constructive* algorithms initially assume a simple network and add nodes and links as warranted [2-8], while *destructive* methods start with a large network and prune off superfluous components [9-12]. Though these algorithms address the problem of topology acquisition, they do so in a highly constrained manner. Because they monotonically modify network structure, constructive and destructive methods limit the traversal of the available architectures in that once an architecture has been explored and determined to be insufficient, a new architecture is adopted, and the old becomes topologically unreachable. Also, these methods often use only a single predefined structural modification, such as "add a fully connected hidden unit," to generate successive topologies. This is a form of *structural hill climbing*, which is susceptible to becoming trapped at structural local minima. In addition, constructive and destructive algorithms make simplifying architectural assumptions to facilitate network induction. For example, Ash [2] allows only feedforward networks; Fahlman [6] assumes a restricted form of recurrence, and Chen et al. [7] explore only fully connected topologies. This creates a situation in which the task is forced into the architecture rather than the architecture being fit to the task.

These deficiencies of constructive and destructive methods stem from inadequate methods for assigning credit to structural components of a network. As a result, the heuristics used are overly-constrained to increase the likelihood of finding any topology to solve the problem. Ideally, the constraints for such a solution should come from the task rather than be implicit in the algorithm.

This paper presents GNARL, a network induction algorithm that simultaneously acquires both network topology and weight values while making minimal architectural restrictions and avoiding structural hill climbing. The algorithm, described in section 3, is an instance of evolutionary programming [13, 14], a class of evolutionary computation that has been shown to perform well at function optimization. Section 2 argues that this class of evolutionary computation is better suited for evolving neural networks than genetic algorithms [15, 16], a more popular class of evolutionary computation. Finally, section 4 demonstrates GNARL's ability to create recurrent networks for a variety of problems of interest.

2.0 Evolving Connectionist Networks

Evolutionary computation provides a promising collection of algorithms for structural and parametric learning of recurrent networks [17]. These algorithms are distinguished by their reliance on a *population* of search space positions, rather than a single position, to locate extrema of a function defined over the search space. During one search cycle, or *generation*, the members of the population are ranked according to a *fitness function*, and those with higher fitness are probabilistically selected to become *parents* in the next generation. New population members, called *offspring*, are created using specialized *reproduction heuristics*. Using the population, reproduction heuristics, and fitness function, evolutionary computation implements a nonmonotonic search that performs well in complex multimodal environments. Classes of evolutionary computation can be distinguished by examining the specific reproduction heuristics employed.

Genetic algorithms (GAs) [15, 16] are a popular form of evolutionary computation that rely chiefly on the reproduction heuristic of *crossover*.[1] This operator forms offspring by recombining representational components from two members of the population without regard to content. This purely structural approach to creating novel population members assumes that components of all parent representations may be freely exchanged without inhibiting the search process.

Various combinations of GAs and connectionist networks have been investigated. Much research concentrates on the acquisition of parameters for a fixed network architecture (e.g., [18 - 21]). Other work allows a variable topology, but disassociates structure acquisition from acquisition of weight values by interweaving a GA search for network topology with a traditional parametric training algorithm (e.g., backpropagation) over weights (e.g., [22, 23]). Some studies attempt to coevolve both the topology and weight values within the GA framework, but as in the connectionist systems described above, the network architectures are restricted (e.g., [24 - 26]). In spite of this collection of studies, current theory from both genetic algorithms and connectionism suggests that GAs are not well-suited for evolving networks. In the following section, the reasons for this mismatch are explored.

---

1. Genetic algorithms also employ other operators to manipulate the population, including a form of mutation, but their distinguishing feature is a heavy reliance on crossover.
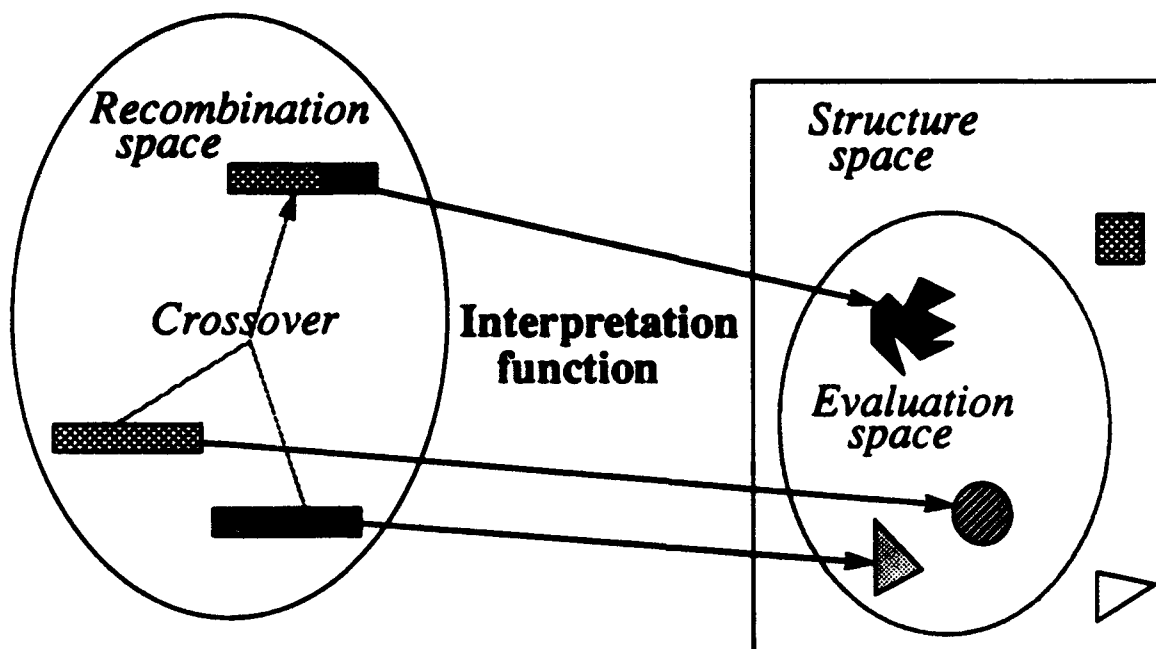
Figure 1. The dual representation scheme used in genetic algorithms. The interpretation function maps between the elements in recombination space on which the search is performed and the subset of structures that can be evaluated as potential task solutions.

2.1 Evolving Networks with Genetic Algorithms

Genetic algorithms create new individuals by recombining the representational components of two member of the population. Because of this commitment to structural recombination, GAs typically rely on two distinct representational spaces (Figure 1). *Recombination space*, usually defined over a set of fixed-length binary strings, is the set of structures to which the genetic operators are applied. It is here that the search actually occurs. *Evaluation space*, typically involving a problem-dependent representation, is the set of structures whose ability to perform a task is evaluated. In the case of using GAs to evolve networks, evaluation space is comprised of a set of networks. An *interpretation function* maps between these two representational spaces. Any set of finite-length bit strings cannot represent all possible networks, thus the evaluation space is restricted to a predetermined set of networks. By design, the dual representation scheme allows the GA to crossover the bit strings without any knowledge of their interpretation as networks. The implicit assumption is that the interpretation function is defined so that the bit strings created by the dynamics of the GA will map to successively better networks.

The dual representation of GAs is an important feature for searching in certain environments. For instance, when it is unclear how to search the evaluation space directly, and when there exists an interpretation function such that searching the space of bit strings by crossover leads to good points in evaluation space, then the dual representation is ideal. It is unclear, however, that there exists an interpretation function that makes dual representation beneficial for evolving neural networks. Clearly, the choice of interpretation function introduces a strong bias into the search, typically by excluding many potentially interesting and useful networks (another example of forcing
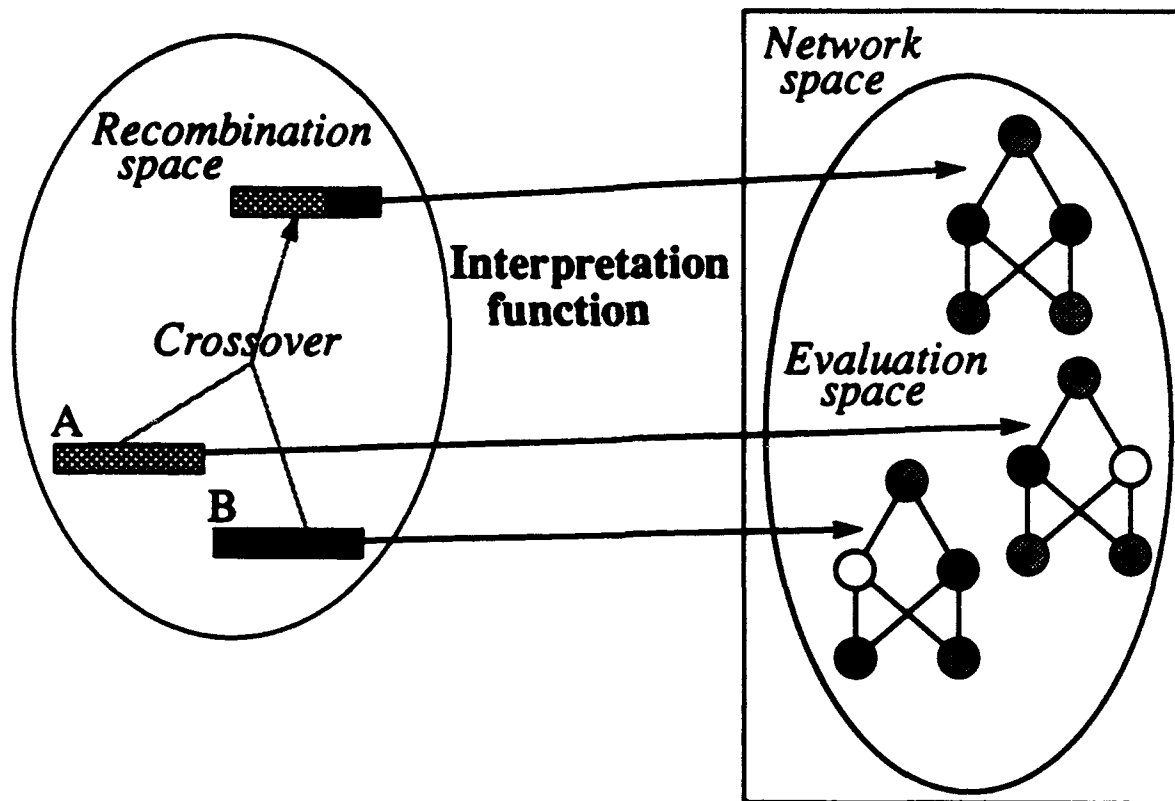
Figure 2. *The competing conventions problem [29]. Bit strings A and B map to structurally and computationally equivalent networks that assign the hidden units in different orders. Because the bit strings are distinct, crossover is likely to produce an offspring that contains multiple copies of the same hidden node, yielding a network with less computational ability than either parent.*

the task into an architecture). Moreover, the benefits of having a dual representation hinge on crossover being an appropriate evolutionary operator for the task for some particular interpretation function; otherwise, the need to translate between dual representations is an unnecessary complication.

Characterizing tasks for which crossover is a beneficial operator is an open question. Current theory suggests that crossover will tend to recombine short, connected substrings of the bit string representation that correspond to above-average task solutions when evaluated [16, 15]. These substrings are called *building blocks*, making explicit the intuition that larger structures with high fitness are built out of smaller structures with moderate fitness. Crossover tends to be most effective in environments where the fitness of a member of the population is reasonably correlated with the expected ability of its representational components [27]. Environments where this is not true are called *deceptive* [28].

There are three forms of deception when using crossover to evolve connectionist networks. The first involves networks that share both a common topology and common weights. Because the interpretation function may be many-to-one, two such networks need not have the same bit string representation (see Figure 2). Crossover will then tend to create offspring that contain repeated components, and lose the computational ability of some of the parents' hidden units. The

---

resulting networks will tend to perform worse than their parents because they do not possess key computational components for the task. Schaffer et al. [29] term this the *competing conventions* problem, and point out that the number of competing conventions grows exponentially with the number of hidden units.

The second form of deception involves two networks with identical topologies but different weights. It is well known that for a given task, a single connectionist topology affords multiple solutions for a task, each implemented by a unique *distributed representation* spread across the hidden units [30, 31]. While the removal of a small number of nodes has been shown to effect only minor alterations in the performance of a trained network [30, 31], the computational role each node plays in the overall representation of the task solution is determined purely by the presence and strengths of its interconnections. Furthermore, there need be no correlation between distinct distributed representations over a particular network architecture for a given task. This seriously reduces the chance that an arbitrary crossover operation between distinct distributed representations will construct viable offspring *regardless of the interpretation function used.*

Finally, deception can occur when the parents differ topologically. The types of distributed representations that can develop in a network vary widely with the number of hidden units and the network's connectivity. Thus, the distributed representations of topologically distinct networks have a greater chance of being incompatible parents. This further reduces the likelihood that crossover will produce good offspring.

In short, for crossover to be a viable operator when evolving networks, the interpretation function must somehow compensate for all the types of deceptiveness described above. This suggests that the complexity of an appropriate interpretation function will more than rival the complexity of the original learning problem. Thus, the prospect of evolving connectionist networks with crossover appears limited in general, and better results should be expected with reproduction heuristics that respect the uniqueness of the distributed representations. This point has been tacitly validated in the genetic algorithm literature by a trend towards a reduced reliance on binary representations when evolving networks (e.g. [32, 33]). Crossover, however, is still commonplace.

2.2 Networks and Evolutionary Programming

Unlike genetic algorithms, evolutionary programming (EP) [14,34] defines representation-dependent mutation operators that create offspring within a specific locus of the parent (see Figure 3). EP's commitment to mutation as the sole reproductive operator for searching over a space is preferable when there is no sufficient calculus to guide recombination by crossover, or when separating the search and evaluation spaces does not afford an advantage.

Relatively few previous EP systems have addressed the problem of evolving connectionist networks. Fogel et al. [35] investigate training feedforward networks on some classic connectionist problems. McDonnell and Waagen [36] use EP to evolve the connectivity of feedforward networks with a constant number of hidden units by evolving both a weight matrix and a connectivity matrix. Fogel [14], [37] uses EP to induce three-layer fully-connected feedforward networks with a variable number of hidden units that employ good strategies for playing Tic-Tac-Toe.
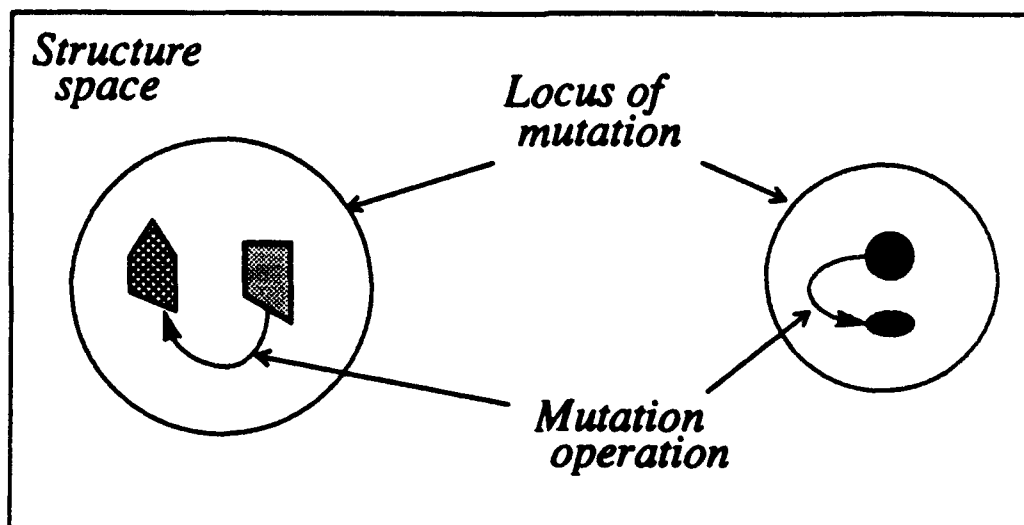
Figure 3. *The evolutionary programming approach to modeling evolution. Unlike genetic algorithms, evolutionary programs perform search in the space of networks. Offspring created by mutation remain within a locus of similarity to their parents.*

In each of the above studies, the mutation operator alters the parameters of network η by the function:

$$w = w + N(0, \alpha\varepsilon(\eta)) \qquad \forall w \in \eta \tag{EQ 1}$$

where w is a weight, $\varepsilon(\eta)$ is the error of the network on the task (typically the mean squared error), α is a user-defined proportionality constant, and $N(\mu, \sigma^2)$ is a gaussian variable with mean μ and variance σ^2. The implementations of structural mutations in these studies differ somewhat. McDonnell and Waagen [36] randomly select a set of weights and alters their values with a probability based on the variance of the incident nodes' activation over the training set; connections from nodes with a high variance having less of a chance of being altered. The structural mutation used in [14, 37] adds or deletes a single hidden unit with equal probability

Evolutionary programming provides distinct advantages over genetic algorithms when evolving networks. First, EP manipulates networks directly, thus obviating the need for a dual representation and the associated interpretation function. Second, by avoiding crossover between networks in creating offspring, the individuality of each network's distributed representation is respected. For these reasons, evolutionary programming provides a more appropriate framework for simultaneous structural and parametric learning in recurrent networks. The GNARL algorithm, presented in the next section and investigated in the remainder of this paper, describes one such approach.

3.0 The GNARL Algorithm

GNARL, which stands for *GeNeralized Acquisition of Recurrent Links*, is an evolutionary algorithm that nonmonotonically constructs recurrent networks to solve a given task. The name

Bias

$m_{in} + b$
input
nodes

m_{out}
output
nodes

at most h_{max}
hidden nodes

Figure 4. Sample initial network. The number of input nodes (m_{in}) and number of output nodes (m_{out}) is fixed for a given task. The presence of a bias node ($b = 0$ or 1) as well as the maximum number of hidden units (h_{max}) is set by the user. The initial connectivity is chosen randomly (see text). The disconnected hidden node does not affect this particular network's computation, but is available as a resource for structural mutations.

GNARL reflects the types of networks that arise from a generalized network induction algorithm performing both structural and parametric learning. Instead of having uniform or symmetric topologies, the resulting networks have "gnarled" interconnections of hidden units which more accurately reflect constraints inherent in the task.

The general architecture of a GNARL network is straightforward. The input and output nodes are considered to be provided by the task and are immutable by the algorithm; thus each network for a given task always has m_{in} input nodes and m_{out} output nodes. The number of hidden nodes varies from 0 to a user-supplied maximum h_{max}. Bias is optional; if provided in an experiment, it is implemented as an additional input node with constant value one. All non-input nodes employ the standard sigmoid activation function. Links use real-valued weights, and must obey three restrictions:

R_1: There can be no links *to* an input node.

R_2: There can be no links *from* an output node.

R_3: Given two nodes x and y, there is at most one link from x to y.

Thus GNARL networks may have no connections, sparse connections, or full connectivity. Consequently, GNARL's search space is:

$S = \{\eta:$ η is a network with real-valued weights,

η satisfies R_1-R_3,

η has $m_{in} + b$ input nodes, where b=1 if a bias node is provided, and 0 otherwise,

η has m_{out} output nodes,

η has i hidden nodes, $0 \leq i \leq h_{max}\}$

R_1-R_3 are strictly implementational constraints. Nothing in the algorithm described below hinges on S being pruned by these restrictions.

3.1 Selection, Reproduction and Mutation of Networks

GNARL initializes the population with randomly generated networks (see Figure 4). The number of hidden nodes for each network is chosen from a uniform distribution over a user-sup-

plied range. The number of initial links is chosen similarly from a second user-supplied range. The incident nodes for each link are chosen in accordance with the structural mutations described below. Once a topology has been chosen, all links are assigned random weights, selected uniformly from the range [-1, 1]. There is nothing in this initialization procedure that forces a node to have *any* incident links, let alone for a path to exist between the input and output nodes. In the experiments below, the number of hidden units for a network in the initial population was selected uniformly between one and five and the number of initial links varied uniformly between one and 10.

In each generation of search, the networks are first evaluated by a user-supplied fitness function $f: S \rightarrow R$, where R represents the reals. Networks scoring in the top 50% are designated as the *parents* of the next generation; all other networks are discarded. This selection method is used in many EP algorithms although competitive methods of selection have also been investigated [14].

Generating an offspring involves three steps: copying the parent, determining the severity of the mutations to be performed, and finally mutating the copy. Network mutations are separated into two classes, corresponding with the types of learning discussed in [1]. *Parametric mutations* alter the value of parameters (link weights) currently in the network, whereas *structural mutations* alter the number of hidden nodes and the presence of links in the network, thus altering the space of parameters.

3.1.1 Severity of Mutations

The severity of a mutation to a given parent, η, is dictated by that network's temperature, $T(\eta)$:

$$T(\eta) = 1 - \frac{f(\eta)}{f_{max}}$$

(EQ 2)

where f_{max} is the maximum fitness for a given task. Thus, the temperature of a network is determined by how close the network is to being a solution for the task. This measure of the network's performance is used to anneal the structural and parametric similarity between parent and offspring, so that networks with a high temperature are mutated severely, and those with a low temperature are mutated only slightly (cf. [38]). This allows a coarse-grained search initially, and a progressively finer-grained search as a network approaches a solution to the task, a process described more concretely below.

3.1.2 Parametric Mutation of Networks

Parametric mutations are accomplished by perturbing each weight w of a network η with gaussian noise, a method motivated by [37, 14]. In that body of work, weights are modified as follows:

$$w = w + N(0, \alpha T(\eta)) \qquad \forall w \in \eta$$

(EQ 3)

where α is a user-defined proportionality constant, and $N(\mu, \sigma^2)$ is a gaussian random variable as before. While large parametric mutations are occasionally necessary to avoid parametric local minima during search, it is more likely they will adversely affect the offspring's ability to perform better than its parent. To compensate, GNARL updates weights using a variant of equation 3. First, the *instantaneous temperature* $\hat{T}$ of the network is computed:

$$\hat{T}(\eta) = U(0, 1) T(\eta) \qquad \text{(EQ 4)}$$

where $U(0, 1)$ is a uniform random variable over the interval $[0, 1]$. This new temperature, varying from 0 to $T(\eta)$, is then substituted into equation 3:

$$w = w + N(0, \alpha \hat{T}(\eta)) \qquad \forall w \in \eta \qquad \text{(EQ 5)}$$

In essence, this modification lessens the frequency of large parametric mutations without disallowing them completely. In the experiments described below, α is one.

3.1.3 Structural Mutation of Networks

The structural mutations used by GNARL alter the number of hidden nodes and the connectivity between all nodes, subject to restrictions R_1-R_3 discussed earlier. To avoid radical jumps in fitness from parent to offspring, structural mutations attempt to preserve the behavior of a network. For instance, new links are initialized with zero weight, leaving the behavior of the modified network unchanged. Similarly, hidden units are added to the network without any incident connections. Links must be added by future structural mutations to determine how to incorporate the new computational unit. Unfortunately, achieving this behavioral continuity between parent and child is not so simple when removing a hidden node or link. Consequently, the deletion of a node involves the complete removal of the node and all incident links with no further modification to compensate for the behavioral change. Similarly, deleting a link removes that parameter from the network.

The selection of which node to remove is uniform over the collection of hidden nodes. Addition or deletion of a link is slightly more complicated in that a parameter identifies the likelihood that the link will originate from an input node or terminate at an output node. Once the class of incident node is determined, an actual node is chosen uniformly from the class. Biasing the link selection process in this way is necessary when there is a large differential between the number of hidden nodes and the number of input or output nodes. This parameter was set to 0.2 in the experiments described in the next section.

Research in [14] and [37] uses the heuristic of adding or deleting at most a single fully connected node per structural mutation. Therefore, it is possible for this method is to become trapped at a structural local minima, although this is less probable than in nonevolutionary algorithms given that several topologies may be present in the population. In order to more effectively search the range of network architectures, GNARL uses a severity of mutation for each separate structural mutation. A unique user-defined interval specifying a range of modification is associated with each of the four structural mutations. Given an interval of $[\Delta_{min}, \Delta_{max}]$ for a particular structural mutation, the number of modifications of this type made to an offspring is given by:

$$\Delta_{min} + \lfloor U[0, 1] \, \hat{r}(\eta) \, (\Delta_{max} - \Delta_{min}) \rfloor \qquad \text{(EQ 6)}$$

Thus the number of modifications varies uniformly over a shrinking interval based on the parent network's fitness. In the experiments below, the maximum number of nodes added or deleted was three while the maximum number of links added or deleted was five. The minimum number for each interval was always one.

3.2 Fitness of a Network

In evolving networks to perform a task, GNARL does not require an explicit target vector – all that is needed is the feedback given by the fitness function f. But if such a vector is present, as in supervised learning, there are many ways of transforming it into a measure of fitness. For example, given a training set $\{(x_1, y_1), (x_2, y_2), ...\}$, three possible measures of fitness for a network η are sum of square errors (equation 7), sum of absolute errors (equation 8), and sum of exponential absolute errors (equation 9):

$$\sum_i (y_i - Out(\eta, x_i))^2 \qquad \text{(EQ 7)}$$

$$\sum_i |y_i - Out(\eta, x_i)| \qquad \text{(EQ 8)}$$

$$\sum_i e^{|y_i - Out(\eta, x_i)|} \qquad \text{(EQ 9)}$$

Furthermore, because GNARL explores the space of networks by mutation and selection, the choice of fitness function does not alter the mechanics of the algorithm. To show GNARL's flexibility, each of these fitness functions will be demonstrated in the experiments below.

4.0 Experiments

In this section, GNARL is applied to several problems of interest. The goal in this section is to demonstrate the abilities of the algorithm on problems from language induction to search and collection. The various parameter values for the program are set as described above unless otherwise noted.

4.1 Williams' Trigger Problem

As an initial test, GNARL induced a solution for the *enable-trigger* task proposed in [39]. Consider the finite state generator shown in Figure 5. At each time step the system receives two input bits, (a, b), representing "enable" and "trigger" signals, respectively. This system begins in state S_1, and switches to state S_2 only when enabled by a=1. The system remains in S_2 until it is triggered by b=1, at which point it outputs 1 and resets the state to S_1. So, for instance, on an input stream $\{(0, 0), (0, 1), (1, 1), (0, 1)\}$, the system will output $\{0, 0, 0, 1\}$ and end in S_1. This simple problem allows an indefinite amount of time to pass between the enable and the trigger inputs;


Figure 5. An FSA that defines the enable-trigger task [39]. The system is given a data stream of bit pairs {(a_1, b_1), (a_2, b_2), ...}, and produces an output of 0's and 1's. To capture this system's input/output behavior, a connectionist network must learn to store state indefinitely.

thus no finite length sample of the output stream will indicate the current state of the system. This forces GNARL to develop networks that can preserve state information indefinitely.

The fitness function used in this experiment was the sum of exponential absolute errors (equation 9). Population size was 50 networks with the maximum number of hidden units restricted to six. A bias node was provided in each network in this initial experiment, ensuring that an activation value of 1 was always available. Note that this does not imply that each node had a nonzero bias; links to the bias node had to be acquired by structural mutation.

Training began with all two input strings of length two, shown in Table 1. After 118 generations (3000 network evaluations[2]), GNARL evolved a network which solved this task for the strings in Table 1 within tolerance of 0.3 on the output units. The training set was then increased to include all 64 input strings of length three and evolution of the networks was allowed to continue. After an additional 422 generations, GNARL once again found a suitable network. At this point, the difficulty of the task was increased a final time by training on all 256 strings of length four. After another 225 generations (~20000 network evaluations total) GNARL once again found a network to solve this task, shown in Figure 6b. Note that there are two completely isolated nodes. Given the fitness function used in this experiment, the two isolated nodes do not effect the network's viability. To investigate the generalization of this network, it was tested over all 4096 unique strings of length six. The outputs were rounded off to the nearest integer, testing only the network's separation of the strings. The network performed correctly on 99.5% of this novel set, generating incorrect responses for only 20 strings.

Figure 7 shows the connectivity of the population member with the best fitness for each generation over the course of the run. Initially, the best network is sparsely-connected and remains sparsely-connected throughout most of the run. At about generation 400, the size and connectivity

---

2. Number of networks evaluated = |population| + generations * |population| * 50% of the population removed each generation, giving 50 + 118 * 50 * 0.5 = 3000 network evaluations for this trial.

Figure 6. *Connectivity of two recurrent networks found in the enable-trigger experiment. (a) The best network of generation 1. (b) The best network of generation 765. This network solves the task for all strings of length eight.*

increases dramatically only to be overtaken by the relatively sparse architecture shown in Figure 6b on the final generation. Apparently, this more sparsely connected network evolved more quickly than the more full architectures that were best in earlier generations. The oscillations between different network architectures throughout the run reflects the development of such competing architectures in the population.

4.2 Inducing Regular Languages

A current topic of research in the connectionist community is the induction of finite state automata (FSAs) by networks with second-order recurrent connections. For instance, Pollack [40] trains sequential cascaded networks (SCNs) over a test set of languages, provided in [41] and

Input	Target Output
{(0, 0), (0, 0)}	{0, 0}
{(0, 0), (0, 1)}	{0, 0}
{(0, 0), (1, 0)}	{0, 0}
{(0, 0), (1, 1)}	{0, 0}
{(0, 1), (0, 0)}	{0, 0}
{(0, 1), (0, 1)}	{0, 0}
{(0, 1), (1, 0)}	{0, 0}
{(0, 1), (1, 1)}	{0, 0}

Input	Target Output
{(1, 0), (0, 0)}	{0, 0}
{(1, 0), (0, 1)}	{0, 1}
{(1, 0), (1, 0)}	{0, 0}
{(1, 0), (1, 1)}	{0, 1}
{(1, 1), (0, 0)}	{0, 0}
{(1, 1), (0, 1)}	{0, 1}
{(1, 1), (1, 0)}	{0, 0}
{(1, 1), (1, 1)}	{0, 1}

**Table 1.** *Initial training data for enable-trigger task.*

Figure 7. *Different network topologies explored by GNARL during the first 540 generations on the enable-trigger problem. The presence of a link between node i and j at generation g is indicated by a dot at position (g, 10 * i + j) in the graph. Note that because node 3 is the output node, there are no connections from it throughout the run. The arrow designates the point of transition between the first two training sets.*

shown in Table 2, using a variation of backpropagation. An interesting result of this work is that the number of states used by the network to implement finite state behavior is potentially infinite. Other studies using the training sets in [41] have investigated various network architectures and training methods, as well as algorithms for extracting FSAs from the trained architectures [42 - 45].

An explicit collection of positive and negative examples, shown in Table 3, that pose specific difficulties for inducing the intended languages is offered in [41]. Notice that the training sets are unbalanced, incomplete and vary widely in their ability to strictly define the intended regular language. GNARL's ability to learn and generalize from these training sets was compared against the training results reported for the second-order architecture used in [42]. Notice that all the languages in Table 2 require recurrent network connections in order to induce the language completely. The type of recurrence needed for each language varies widely. For instance, languages 1 through 4 require an incorrect input be remembered indefinitely, forcing the network to develop an analog version of a trap state. Networks for language 6, however, must parse and count indi-

Language	Description
1	1^*
2	$(1\,0)^*$
3	no odd length 0 strings anytime after an odd length 1 string
4	no more than two 0s in a row
5	an even sum of 10s and 01s, pairwise
6	(number of 1s - number of 0s) mod 3 = 0
7	$0^*1^*0^*1^*$

Table 2. Regular languages to be induced.

vidual inputs, potentially changing state from accept to reject or vice versa on each successive input.

The results obtained in [42] are summarized in Table 4. The table shows the number of networks evaluated to learn the training set and the accuracy of generalization for the learned network to the intended regular language. Accuracy is measured as the percentage of strings of

Language	Positive Instances	Negative Instances
1	ε, 1, 11, 111, 1111, 11111, 111111, 1111111, 11111111	0, 10, 01, 00, 011, 110, 000, 11111110, 10111111
2	ε, 10, 1010, 101010, 10101010, 10101010101010	1, 0, 11, 00, 01, 101, 100, 1001010, 10110, 110101010
3	ε, 1, 0, 01, 11, 00, 100, 110, 111, 000, 100100, 110000011100001, 111101100010011100	10, 101, 010, 1010, 110, 1011, 10001, 111010, 1001000, 11111000, 0111001101, 11011100110
4	ε, 1, 0, 10, 01, 00, 100100, 001111110100, 0100100100, 11100, 010	000, 11000, 0001, 000000000, 00000, 0000, 11111000011, 1101010000010111, 1010010001
5	ε, 11, 00, 001, 0101, 1010, 1000111101, 1001100001111010, 111111, 0000	1, 0, 111, 010, 000000000, 1000, 01, 10, 1110010100, 010111111110, 0001, 011
6	ε, 10, 01, 1100, 101010, 111, 000000, 0111101111, 100100100	1, 0, 11, 00, 101, 011, 11001, 1111, 00000000, 010111, 10111101111, 1001001001
7	ε, 1, 0, 10, 01, 11111, 000, 00110011, 0101, 0000100001111, 00100, 011111011111, 00	1010, 00110011000, 0101010101, 1011010, 10101, 010100, 101001, 100100110101

Table 3. Training sets for the languages of Table 2 from [41].

Language	Average evaluations	Average % accuracy	Fewest evaluations	Best % accuracy
1	3033.8	88.98	28	100.0
2	4522.6	91.18	807	100.0
3	12326.8	64.87	442	78.31
4	4393.2	42.50	60	60.92
5	1587.2	44.94	368	66.83
6	2137.6	23.19	306	46.21
7	2969.0	36.97	373	55.74

Table 4. Speed and generalization results reported by [42] for learning the data sets of Table 3.

length 10 or less that are correctly classified by the network. For comparison, the table lists both the average and best performance of the five runs reported in [42].

This experiment used a population of 50 networks, each limited to at most eight hidden units. Each run lasted at most 1000 generations, allowing a maximum of 25050 networks to be evaluated for a single data set. Two experiments were run for each data set, one using the sum of absolute errors (SAE) and the other using sum of square errors (SSE). The error for a particular string was computed only for the final output of the network after the entire string plus three trailing "null" symbols had been entered, one input per time step. The concatenation of the trailing null symbols was used to identify the end of the string and allow input of the null string, a method also used in [42]. Each network had a single input and output and no bias node was provided. The three possible logical inputs for this task, 0, 1, and null, were represented by activations of -1, 1, and 0, respectively. The tolerance for the output value was 0.1, as in [42].

Table 5 shows for both fitness functions the number of evaluations until convergence and the accuracy of the best evolved network. Only four of the runs, each of those denoted by a '+' in the table, failed to produce a network with the specified tolerance in the allotted 1000 generations. In the runs using SAE, the two runs that did not converge had not separated a few elements of the associated training set and appeared to be far from discovering a network that could correctly classify the complete training set. Both of the uncompleted runs using SSE successfully separated the data sets but had not done so to the 0.1 tolerance within the 1000 generation limit. Figure 8 compares the number of evaluations by GNARL to the average number of evaluations reported in [42]. As the graph shows, GNARL consistently evaluates more networks, but not a disproportionate number. Considering that the space of networks being searched by GNARL is much larger than the space being searched by [42], these numbers appear to be within a tolerable increase.

The graph of Figure 9 compares the accuracy of the GNARL networks to the *average* accuracy found in [42] over five runs. The GNARL networks consistently exceeded the average accuracy found in [42].

Figure 8. *The number of network evaluations required to learn the seven data sets of Table 3. GNARL (using both SAE and SSE fitness measures) compared to the average number of evaluations for the five runs described in [42].*

These results demonstrate GNARL's ability to simultaneously acquire the topology and weights of recurrent networks, and that this can be done within a comparable number of network evaluations as training a network with static architecture on the same task. GNARL also appears to generalize better consistently, possibly due to its selective inclusion and exclusion of some links.

Language	Evaluations (SAE)	% Accuracy (SAE)	Evaluations (SSE)	% Accuracy (SSE)
1	3975	100.00	5300	99.27
2	5400	96.34	13975	73.33
3	25050[+]	58.87	18650	68.00
4	15775	92.57[*]	21850	57.15
5	25050[+]	49.39	22325	51.25
6	21475	55.59[*]	25050[+]	44.11
7	12200	71.37[*]	25050[+]	31.46

Table 5. *Speed and generalization results for GNARL to train recurrent networks to recognize the data sets of Table 3.*

Figure 9. *Percentage accuracy of evolved networks on languages in Table 2. GNARL (using SAE and SSE fitness measures) compared to average accuracy of the five runs in [42].*

4.3 The Ant Problem

GNARL was tested on a complex search and collection task – the *Tracker* task described in [46], and further investigated in [47]. In this problem, a simulated ant is placed on a two-dimensional toroidal grid that contains a trail of food. The ant traverses the grid, collecting any food it contacts along the way. The goal of the task is to discover an ant which collects the maximum number of pieces of food in a given time period. (Figure 10).

Following [46], each ant is controlled by a network with two input nodes and four output nodes (Figure 11). The first input node denotes the presence of food in the square directly in front of the ant; the second denotes the absence of food in this same square, restricting the possible legal inputs to the network to (1, 0) or (0, 1). Each of the four output units corresponds to a unique action: move forward one step, turn left 90°, turn right 90°, or no-op. At each step, the action whose corresponding output node has maximum activation is performed. As in the original study [46], no-op allows the ant to remain at a fixed position while activation flows along recurrent connections. Fitness is defined as the number of grid positions cleared within 200 time steps. The task is difficult because simple networks can perform surprisingly well; the network shown in Figure 11 collects 42 pieces of food before spinning endlessly at position A (in Figure 10), illustrating a very high local maximum in the search space.

The experiment used a population of 100 networks, each limited to at most nine hidden units, and did not provide a bias node. In the first run (2090 generations), GNARL found a network (Figure 12b) that clears 81 grid positions within the 200 time steps. When this ant is run for an additional 119 time steps, it successfully clears the entire trail. To understand how the network traverses the path of food, consider the simple FSA shown in Figure 13, hand-crafted in [46] as an approximate solution to the problem. This simple machine receives a score of 81 in the allotted

Figure 10. *The ant problem. The trail is connected initially, but becomes progressively more difficult to follow. The underlying 2-d grid is toroidal, so that position "A" is the first break in the trail – it is simple to reach this point. Positions "B" and "C" indicate the only two positions along the trail where the ant discovered in run 1 behaves differently from the 5-state FSA of [46] (see Figure 13).*

200 time steps, and clears the entire trail only five time steps faster than the network in Figure 12b. A step by step comparison indicates there is only a slight difference between the two. GNARL's evolved network follows the general strategy embodied by this FSA at all but two places, marked as positions B and C in Figure 10. Here the evolved network makes a few additional moves, accounting for the slightly longer completion time.


Figure 11. *The semantics of the I/O units for the ant network. The first input node denotes the presence of food in the square directly in front of the ant; the second denotes the absence of food in this same square. This particular network finds 42 pieces of food before spinning endlessly in place at position P, illustrating a very deep local minimum in the search space.*

(a)

(b)

Figure 12. The Tracker Task, first run. (a) The best network in the initial population. Nodes 0 and 1 are input, nodes 5-8 are output, and nodes 2-4 are hidden nodes. (b) Network induced by GNARL after 2090 generations. Forward links are dashed; bidirectional links and loops are solid. The light gray connection between nodes 8 and 13 is the sole backlink. This network clears the trail in 319 epochs.

Figure 14 illustrates the strategy the network uses to implement the FSA by showing the state of the output units of the network over three different sets. Each point is a triple of the form (*move, right, left*).[3] Figure 14a shows the result of supplying to the network 200 "food" inputs – a fixed point that executes "Move." Figure 14b shows the sequence of states reached when 200 "no food" signals are supplied to the network – a collection of points describing a limit cycle of length five that repeatedly executes the sequence "Right, Right, Right, Right, Move." These two attractors determine the response of the network to the task (Figure 14c, d); the additional points in Fig-

---

3. No-op is not shown because it was never used in the final network.

Figure 13. FSA hand-crafted for the Tracker task in [46]. The large arrow indicates the initial state. This simple system implements the strategy "move forward if there is food in front of you, otherwise turn right four times, looking for food. If food is found while turning, pursue it, otherwise, move forward one step and repeat." This FSA traverses the entire trail in 314 steps, and gets a score of 81 in the allotted 200 time steps.

ure 14c are transients encountered as the network alternates between these attractors. The differences in the number of steps required to clear the trail between the FSA of Figure 13 and GNARL's network arise due to the state of the hidden units when transferring from the "food" attractor to the "no food" attractor.

However, not all evolved network behaviors are so simple as to approximate an FSA [40]. In a second run (1595 generations) GNARL induced a network that cleared 82 grid points within the 200 time steps. Figure 15 demonstrates the behavior of this network. Once again, the "food" attractor, shown in Figure 15a, is a single point in the space that always executes "Move." The "no food" behavior, however, is not an FSA; instead, it is a quasiperiodic trajectory of points shaped like a "D" in output space (Figure 15b). The placement of the "D" is in the "Move / Right" corner of the space and encodes a complex alternation between these two operations (see Figure 15d).

In contrast, research in [46] uses a genetic algorithm on a population of 65,536 bit strings with a direct encoding to evolve only the weights of a neural network with five hidden units to solve this task. The particular network architecture in [46] uses Boolean threshold logic for the hidden units and an identity activation function for the output units. The first GNARL network was discovered after evaluating a total of 104,600 networks while the second was found after evaluating 79,850. The experiment reported in [46] discovered a comparable network after about 17 generations. Given [46] used a population size of 65,536 and replaced 95% of the population each generation, the total number of network evaluations to acquire the equivalent network was 1,123,942. This is 10.74 and 14.07 times the number of networks evaluated by GNARL in the two runs. In spite of the differences between the two studies, this significant reduction in the number of evaluations provides empirical evidence that crossover may not be best suited to the evolution of networks.

Figure 14. Limit behavior of the network that clears the trail in 319 steps. Graphs show the state of the output units Move, Right, Left. (a) Fixed point attractor that results for sequence of 500 "food" signals; (b) Limit cycle attractor that results when a sequence of 500 "no food" signals is given to network; (c) All states visited while traversing the trail; (d) The path of the ant on an empty grid. The Z axis represents time. Note that x is fixed, and y increases monotonically at a fixed rate. The large jumps in y position are artifacts of the toroidal grid.

5.0 Conclusions

Allowing the task to specify an appropriate architecture for its solution should, in principle, be the defining aspect of the complete network induction problem. By restricting the space of networks explored, constructive, destructive, and genetic algorithms only partially address the problem of topology acquisition. GNARL's architectural constraints R_1-R_3 similarly reduce the search space, but to a far less degree. Furthermore, none of these constraints is necessary, and their

Figure 15. Limit behavior of the network of the second run. Graphs show the state of the output units Move, Right, Left. (a) Fixed point attractor that results for sequence of 3500 "food" signals; (b) Limit cycle attractor that results when a sequence of 3500 "no food" signals is given to network; (c) All states visited while traversing the trail; (d) The path of the ant on an empty grid. The z axis represents time. The ant's path is comprised of a set of "railroad tracks." Along each track, tick marks represent back and forth movement. At the junctures between tracks, a more complicated movement occurs. There are no artifacts of the toroidal grid in this plot, all are actual movements (cf. Figure 14d).

removal would affect only ease of implementation. In fact, no assumed features of GNARL's networks are essential for the algorithm's operation. GNARL could even use nondifferentiable activation functions, a constraint for backpropagation.

GNARL's minimal representational constraints would be meaningless if not complemented by appropriate search dynamics to traverse the space of networks. First, unlike constructive and destructive algorithms, GNARL permits a nonmonotonic search over the space of network topol-

ogies. Consider that in monotonic search algorithms, the questions of *when* and *how* to modify structure take on great significance because a premature topological change cannot be undone. In contrast, GNARL can revisit a particular architecture at any point, but for the architecture to be propagated it must confer an advantage over other competing topologies. Such a non-linear traversal of the space is imperative for acquiring appropriate solutions because the efficacy of the various architectures changes as the parametric values are modified.

GNARL allows multiple structural manipulations to a network within a single mutation. As discussed earlier, constructive and destructive algorithms define a unit of modification, e.g., "add a fully connected hidden node." Because such singular structural modifications create a "one-unit structural horizon" beyond which no information is available, such algorithms may easily fixate on an architecture that is better than networks one modification step away, but worse than those two or more steps distant. In GNARL, several nodes and links can be added or deleted with each mutation, the range being determined by user-specified limits and the current ability of the network. This simultaneous modification of the structural and parametric modifications based on fitness allows the algorithm to discover appropriate networks quickly especially in comparison to evolutionary techniques that do not respect the uniqueness of distributed representations.

Finally, as in all evolutionary computation, GNARL maintains a population of structures during the search. This allows the algorithm to investigate several differing architectures in parallel while avoiding over-commitment to a particular network topology.

These search dynamics, combined with GNARL's minimal representational constraints make the algorithm extremely versatile. Of course, if topological constraints are known a priori, they should be incorporated into the search. *But these should be introduced as part of the task specification rather than being built into the search algorithm.* Because the only requirement on a fitness function f is that $f: S \rightarrow \mathcal{R}$, diverse criteria can be used to rate a network's performance. For instance, the first two experiments described above evaluated networks based on a desired input/output mapping; the Tracker task experiment, however, considered overall network performance, not specific mappings. Other criteria could also be introduced, including specific structural constraints (e.g., minimal number of hidden units or links) as well as constraints on generalization. In some cases, strong task restrictions can even be implicit in simple fitness functions [48].

The dynamics of the algorithms guided by the task constraints represented in the fitness function allow GNARL to empirically determine an appropriate architecture. Over time, the continual cycle of *test-prune-reproduce* will constrain the population to only those architectures that have acquired the task most rapidly. Inappropriate networks will not be indefinitely competitive and will be removed from the population eventually.

Complete network induction must be approached with respect to the complex interaction between network topology, parametric values, and task performance. By fixing topology, gradient descent methods can be used to discover appropriate solutions. But the relationship between network structure and task performance is not well understood, and there is no "backpropagation" through the space of network architectures. Instead, the network induction problem is approached with heuristics that, as described above, often restrict the available architectures, the dynamics of the search mechanism, or both. Artificial architectural constraints (such as "feedforwardness") or overly constrained search mechanisms can impede the induction of entire classes of behaviors,

while forced structural liberties (such as assumed full recurrence) may unnecessarily increase structural complexity or learning time. By relying on a simple stochastic process, GNARL strikes a middle ground between these two extremes, allowing the network's complexity and behavior to emerge in response to the requirements of the task.

6.0 Acknowledgments

7.0 References

[1]   A. G. Barto. Connectionist learning for control. In W. T. Miller III, R. S. Sutton, and P. J. Werbos, editors, *Neural Networks for Control*, chapter 1, pages 5–58. MIT Press, Cambridge, 1990.

[2]   T. Ash. Dynamic node creation in backpropagation networks. *Connection Science*, 1(4):365–375, 1989.

[3]   M. Frean. The upstart algorithm: A method for constructing and training feed-forward neural networks. Technical Report Preprint 89/469, Edinburgh Physics Dept, 1990.

[4]   S. J. Hanson. Meiosis networks. In D. Touretzky, editor, *Advances in Neural Information Processing Systems 2*, pages 533–541. Morgan Kaufmann, San Mateo, CA, 1990.

[5]   S. E. Fahlman and C. Lebiere. The cascade-correlation architecture. In D. S. Touretsky, editor, *Advances in Neural Information Processing Structures 2*, pages 524–532. Morgan Kaufmann, San Mateo, CA, 1990.

[6]   S. Fahlman. The recurrent cascade-correlation architecture. In R. Lippmann, J. Moody, and D. Touretzky, editors, *Advances in Neural Information Processing Systems 3*, pages 190–196. Morgan Kaufmann, San Mateo, CA, 1991.

[7]   D. Chen, C. Giles, G. Sun, H. Chen, Y. Less, and M. Goudreau. Constructive learning of recurrent neural networks. *IEEE International Conference on Neural Networks*, 3:1196–1201, 1993.

[8]   M. R. Azimi-Sadjadi, S. Sheedvash, and F. O. Trujillo. Recursive dynamic node creation in multilayer neural networks. *IEEE Transactions on Neural Networks*, 4(2):242–256, 1993.

[9]   M. Mozer and P. Smolensky. Skeletonization: A technique for trimming the fat from a network via relevance assessment. In D. Touretzky, editor, *Advances in Neural Information Processing Systems 1*, pages 107–115. Morgan Kaufmann, San Mateo, CA, 1989.

[10] Y. L. Cun, J. Denker, and S. Solla. Optimal brain damage. In D. Touretzky, editor, *Advances in Neural Information Processing Systems 2*. Morgan Kaufmann, San Mateo, CA, 1990.

[11] B. Hassibi and D. G. Stork. Second order derivatives for network pruning: Optimal brain surgeon. In S. J. Hanson, J. D. Cowan, and C. L. Giles, editors, *Advances in Neural Information Processing Systems 5*, pages 164–171. Morgan Kaufmann, San Mateo, CA, 1993.

[12] C. W. Omlin and C. L. Giles. Pruning recurrent neural networks for improved generalization performance. Technical Report Tech Report No 93-6, Computer Science Department, Rensselaer Polytechnic Institute, April 1993.

[13] L. J. Fogel, A. J. Owens, and M. J. Walsh. *Artificial Intelligence through Simulated Evolution*. John Wiley & Sons, New York, 1966.

[14] D. B. Fogel. *Evolving Artificial Intelligence*. Ph.D. thesis, University of California, San Diego, 1992.

[15] J. H. Holland. *Adaptation in Natural and Artificial Systems*. The University of Michigan Press, Ann Arbor, MI, 1975.

[16] D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley Publishing Company, Inc., Reading, MA, 1989.

[17] D. B. Fogel. An introduction to simulated evolutionary optimization. This issue.

[18] A. P. Wieland. Evolving neural network controllers for unstable systems. In *IEEE International Joint Conference on Neural Networks*, pages II-667 – II-673, IEEE Press, Seattle, WA, 1990.

[19] D. Montana and L. Davis. Training feedforward neural networks using genetic algorithms. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 762–767, Morgan Kaufmann, San Mateo, CA, 1989.

[20] D. Whitley, T. Starkweather, and C. Bogart. Genetic algorithms and neural networks: Optimizing connections and connectivity. *Parallel Computing*, 14:347–361, 1990.

[21] R. D. Beer and J. C. Gallagher. Evolving dynamical neural networks for adaptive behavior. *Adaptive Behavior*, 1(1):91–122, 1992.

[22] G. F. Miller, P. M. Todd, and S. U. Hegde. Designing neural networks using genetic algorithms. In J. D. Schaffer, editor, *Proceedings of the Third International Conference on Genetic Algorithms*, pages 379–384. Morgan Kaufmann, San Mateo, CA, 1989.

[23] R. K. Belew, J. McInerney, and N. N. Schraudolf. Evolving networks: Using the genetic algorithm with connectionist learning. Technical Report CS90-174, University of California, San Diego, June 1990.

[24] J. Torreele. Temporal processing with recurrent networks: An evolutionary approach. In R. K. Belew and L. B. Booker, editors, *Fourth International Conference on Genetic Algorithms*, pages 555–561. Morgan Kaufmann, San Mateo, California, 1991.

[25] M. A. Potter. A genetic cascade-correlation learning algorithm. In *Proceedings of COGANN-92 International Workshop on Combinations of Genetic Algorithms and Neural Networks*, 1992.

[26] N. Karunanithi, R. Das, and D. Whitley. Genetic cascade learning for neural networks. In *Proceedings of COGANN-92 International Workshop on Combinations of Genetic Algorithms and Neural Networks*, 1992.

[27] D. E. Goldberg. Genetic algorithms and Walsh functions: Part 2, Deception and its analysis. *Complex Systems*, 3:153–171, 1989.

[28] D. E. Goldberg. Genetic algorithms and Walsh functions: Part 1, A gentle introduction. *Complex Systems*, 3:129–152, 1989.

[29] J. D. Schaffer, D. Whitley, and L. J. Eshelman. Combinations of genetic algorithms and neural networks: A survey of the state of the art. In *Proceedings of COGANN-92 International Workshop on Combinations of Genetic Algorithms and Neural Networks*, 1992.

[30] G. E. Hinton, J. L. McClelland, and D. E. Rumelhart. Distributed representations. In D. E. Rumelhart and J. L. McClelland, editors, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, volume 1: Foundations, pages 77–109. MIT Press, Cambridge, MA, 1986.

[31] T. J. Sejnowski and C. R. Rosenberg. Parallel networks that learn to pronounce english text. *Complex Systems*, 1:145–168, 1987.

[32] J. Koza and J. Rice. Genetic generation of both the weights and architecture for a neural network. In *IEEE International Joint Conference on Neural Networks*, pages II-397 – II-404, Seattle, WA, IEEE Press, 1991.

[33] R. Collins and D. Jefferson. An artificial neural network representation for artificial organisms. In H. P. Schwefel and R. Manner, editors, *Parallel Problem Solving from Nature*. Springer-Verlag, 1991.

[34] D. B. Fogel. A brief history of simulated evolution. In D. B. Fogel and W. Atmar, editors, *Proceedings of the First Annual Conference on Evolutionary Programming*, Evolutionary Programming Society, La Jolla, CA., 1992.

[35] D. B. Fogel, L. J. Fogel, and V. W. Porto. Evolving neural networks. *Biological Cybernetics*, 63:487–493, 1990.

[36] J. R. McDonnell and D. Waagen. Determining neural network connectivity using evolutionary programming. In *Twenty-fifth Asilomar Conferences on Signals, Systems, and Computers*, Monterey, CA, 1992.

[37] D. B. Fogel. Using evolutionary programming to create neural networks that are capable of playing Tic-Tac-Toe. In *International Conference on Neural Networks*, pages 875–880. IEEE Press, San Francisco, CA, 1993.

[38] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220:671–680, 1983.

[39] R. J. Williams. *Adaptive State Representation and Estimation Using Recurrent Connectionist Networks*, chapter 4, pages 97–114. MIT Press, Cambridge, MA, 1990.

[40] J. B. Pollack. The induction of dynamical recognizers. *Machine Learning*, 7:227–252, 1991.

[41] M. Tomita. Dynamic construction of finite automata from examples using hill-climbing. In *Proceedings of the Fourth Annual Conference of the Cognitive Science Society*, pages 105–108, Ann Arbor, MI, 1982.

[42] R. L. Watrous and G. M. Kuhn. Induction of finite-state automata using second-order recurrent networks. In *Advances in Neural Information Processing 4*. Morgan Kaufmann, San Mateo, CA, 1992.

[43] C. L. Giles, G. Z. Sun, H. H. Chen, Y. C. Lee, and D. Chen. Higher order recurrent networks & grammatical inference. In D. S. Touretsky, editor, *Advances in Neural Information Processing Systems 2*, pages 380-387. Morgan Kaufmann, San Mateo, CA, 1990.

[44] C. L. Giles, C. B. Miller, D. Chen, G. Z. Sun, H. H. Chen, and Y. C. Lee. Extracting and learning an unknown grammar with recurrent neural networks. In *Advances in Neural Information Processing 4*. Morgan Kaufmann, San Mateo, CA, 1992.

[45] Z. Zeng, R. M. Goodman, and P. Smyth. Learning finite state machines with self-clustering recurrent networks. *Neural Computation*, to appear.

[46] D. Jefferson, R. Collins, C. Cooper, M. Dyer, M. Flowers, R. Korf, C. Taylor, and A. Wang. Evolution as a theme in artificial life: The genesys/tracker system. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II: Proceedings of the Workshop on Artificial Life*, pages 549-577. Addison-Wesley, 1991.

[47] J. Koza. Genetic evolution and co-evolution of computer programs. In J. D. F. Christopher G. Langton, Charles Taylor and S. Rasmussen, editors, *Artificial Life II*. Addison Wesley Publishing Company, Reading Mass., 1992.

[48] P. J. Angeline and J. B. Pollack. Competitive environments evolve better solutions for complex tasks. In S. Forrest, editor, *Genetic Algorithms: Proceedings of the Fifth International Conference (GA93)*, Morgan Kaufmann, San Mateo, CA, 1993.